Activity-Based Models: 1994-2009

MIT ITS Lab Presentation
March 10, 2009

John L Bowman, Ph.D.
John_L_Bowman@alum.mit.edu
JBowman.net

DAY ACTIVITY SCHEDULE APPROACH

(page 1) In 1994, Moshe and I began developing what I came to call the Activity Schedule Approach to travel demand forecasting.

Since then the approach has been further developed and implemented in several working model systems. Mark Bradley and I have continued to develop the approach each time a planning organization asks us to implement it with them and is interested in enhanced capabilities. Along the way, we have also incorporated features developed with or by others who have also been working in this area.

(page 2) So today, I will be describing what we've been up to for the last ten years, and what I think might lie ahead.

(page 3) The first few slides that I will show you today come from presentations I gave back in the mid 90s.

This simple graphic, which comes from the first TRB presentation I ever gave, shows how Moshe Ben-Akiva, Dinesh Gopinath and I proposed to model a person's day activity pattern that overarches and ties together tours in a nested structure of discrete choice models spanning a day.

Tours are modeled conditional on the activity pattern, with lower priority tours modeled conditional on higher priority tours

And the activity pattern is affected by the ease of carrying out the activities, as measured in the tour models.

(page 4) The day activity schedule can be viewed as an evolutionary step in the use of disaggregate discrete choice methods to model travel demand.

--Trip-based models were developed in the early 70s

--And these were followed, especially in Europe, by tour based models.

The essence of our contribution at the time was to integrate the representation of travel demand across an entire day.

(page 5) Although the exact representation of the activity pattern has varied from time to time, this slide from that first presentation still captures the essence.

--The activity pattern identifies the number and purposes of the tours that a person conducts in their day

--and the number and purposes of additional stops on those tours.

(page 6) For each tour, the timing, destination and travel mode are modeled

--for the tour

--and for each stop on the tour

(page 7) The day activity schedule model operates iteratively with equilibrium assignment models to generate its predictions.

It essentially substitutes for trip generation, distribution and mode split in a traditional 4-step model system.

This makes it manageable for a planning agency using a traditional 4-step model to advance to the activity schedule approach.

(page 8) By the end of 1994 I had implemented the MIT prototype. This slide from my 1995 TRB presentation gave a report on how well we had achieved our objectives.

For the most part we were very pleased with the results

--although we hadn’t yet actually integrated the model with the traffic assignment models

--and we considered our use of five time periods for time-ofday modeling inadequate.

WHAT’S HAPPENING IN THE US?

(page 9) So, what has come of the approach since then?

(page 10) Since then, metropolitan areas have gradually begun to adopt this and similar approaches. And the pace of adoption seems to be picking up. Today there are five metro areas with working model systems, another four in an advanced stage of development,...

(page 11) ...and at least six that are about to get started.

(page 12) After a long period of skepticism and resistance, momentum is growing in the United States for more widespread adoption of these methods. Advocates believe that they are more well-suited for dealing with today’s political issues.

(page 13) TRB Special Report 288 has been an important impetus for this shift (review the bullet points)

(page 14) The report issued a strong condemnation of the current state of the practice (read the bullets)

(page 15) And it made some strong recommendations.... (read the slide)

The federal sponsors have taken these recommendations seriously and have begun to act on them.

And the report seems to have caught the attention of metropolitan agencies.

(page 16) Independent of the federal study and report, activity-based modeling has made its way into the California state law and regulation (read the slides)

So, although nowhere do these documents actually define what activity-based models are, they refer to them explicitly and say that they should be developed and used.

(page 17) And it seems clear that they are referring to the model systems that are being developed and used in these metro areas.

BASICS OF THE "ACTIVITY-BASED" MODEL SYSTEMS

(page 18) So, let's spend a few minutes seeing what these models are like.

(page 19) To do this I start by showing the framework of a trip-based model system, since

--that is what these agencies are familiar with

--AND much of the framework and components are the same

(page 20) The transport model system takes inputs from the urban forecasting and planning process:

(page 21)   --transport networks representing potential future scenarios

(page 22)   --zonal attributes representing forecasts of size and distribution of employment and population in the region

(page 23)   --and socioeconomic attributes of the population.

(page 24) It then predicts zone-to-zone trip flows

(page 25) --and assigns those trips to the network

(page 26) This occurs iteratively to achieve a final result where

--inputs to the trip demand models are consistent with

--predicted link flows and travel speeds coming from assignment of the demand they generate

(page 27) In some cases, accessibility measures from the transport model system serve as input to the land use forecasting models

(page 28) The trip demand includes: --trips associated with special generators, such as airports (page 29) --trips carried out by persons who live outside the region

(page 30) --trips conducted by residents of the region

(page 31) --and commercial movements

(page 32) These trips are usually modeled by separate demand models and the results are combined into trip matrices for highway and transit assignment.

Here is where we can start talking of the difference introduced by an activity-based model….

(page 33) In the existing so-called activity-based model systems, only the trips of residents are modeled using the activitybased model.

(page 34) These are produced by a household travel demand simulator.

--It generates a synthetic population representing the future population of the region,

--and predicts activities and trips for every member of each synthetic household.

(page 35) In some cases, the activity-based model uses a large database of parcel attributes instead of only zonal attributes.

--It takes a lot of resources to generate the parcel database

--But it is very desirable because it provides substantially better information about

--transit accessibility, --travel times for non-motorized and short trips, --and destination attractiveness in heterogeneous zones.

(page 36) The main output of an activity-based travel demand simulator is a detailed itinerary for every person in every household

(page 37) It comes in the form of tables or lists that identify attributes

--of each household

--each person in that household

--each person's day

--each travel tour in their day

--and each trip on that tour

(page 38) These trips are then aggregated and combined with the other predicted trips into matrices for assignment.

So, at this point we see all the components of a so-called activity-based model system needed for transport forecasting.

TECHNICAL EVOLUTION

(page 39) Now I'd like to give you a few highlights from the evolution of these model systems over time.

(page 40) Among all the US projects,

(page 41) the models in Portland, San Francisco, Sacramento, Denver, and now Seattle, are direct descendents of the model I developed with Moshe at MIT. So I'll be looking with you at some of these.

(page 42) Our MIT results caught the attention of Keith Lawton, the director of modeling at Portland Metro, and he decided to implement the approach there.

(page 43) I led a design effort to fill in details missing from the prototype, and Mark Bradley subsequently developed an operational model system for Metro while I continued my studies at MIT.

(page 44) We included models of long term work and school location choice that condition the day activity pattern....

(page 45) ...adding another level to the basic hierarchy of models

(page 46) We also modeled details of the intermediate stops on tours, and integrated the activity schedule model with traffic assignment, which I hadn't done in the prototype.

(page 47) We also included other important details And the model system was used for policy analysis.

(page 48) Here are a couple results from the earliest documented application, reported before the equilibration with traffic assignment was implemented. At the time we were apologetic about the enormous size of the price increases in the test scenarios (changing variable costs from 8 cents per mile to 16)

They showed several important aspects of the model sensitivity:

(page 49) --the decrease in tours shows that activity generation is sensitive to LOS, but less so than VMT

(page 50) --The decrease in tours and miles came primarily from auto travel

(page 51) --Elasticity of demand was greater for discretionary and maintenance than for work purposes, when the costs increased for all times of day (page 52) and when the costs changed only in the peak period, there was a shift from the peak periods to the off-peak periods (right hand column, off-peak row)

(page 53) Unfortunately, the Portland model was mothballed because Metro's money was short and they decided to focus their attention on TranSIMS, where Federal money was available.

But by then San Francisco County was interested in developing an advanced model, and it was decided to implement a version of the activity-schedule approach, with somewhat simplified models and integration. The SFCTA model was developed by PB, Mark Bradley and Cambridge Systematics.

(page 54) It was the first activity-based model to be used extensively and on an ongoing basis for policy analysis and decision-making.

(page 55) And it has been enhanced over time. For example, it has recently been improved to

--include toll vs free alternatives in mode choice

--to use randomly distributed values of time in the models, where values of time are drawn from VOT distributions as part of the simulation.

(page 56) The next major implementation of the approach didn't occur until 2005 in Sacramento. I will spend a bit more time describing the SACOG model, which we call DaySim, since it is the most recently completed new implementation now in use.

(page 57) As with prior implementations, we respecified the day activity pattern somewhat. The question in specifying the pattern model is to decide which aspects of the schedule to include in the pattern model itself, and which to model lower in the hierarchy within individual tours.

(page 58) We decided to identify tour and stop purposes more specifically in the pattern model itself,

--and NOT to determine whether stops occur before or after the primary destination of specific tours.

(page 59) Modeling detailed outcomes was important to SACOG, especially in the spatial dimension.

(page 60) All the models distinguish seven distinct purposes

(page 61) We model the choice of a single 30-minute time period, for every time-of-day outcome in the model

(page 62) And we model all location choices as choice of parcel, instead of the customary traffic analysis zone

(page 63) So, looking at the entire model hierarchy, we see that new detail is being added at many levels in the model system (page 64) In a hierarchical system of models integration is very important, if the component choices are actually correlated

(page 65) We achieve “downward integration” by conditioning lower level models on upper level outcomes.

An important aspect of this is the enforcement of time-space constraints.

(page 66) Whenever a tour or intermediate stop is scheduled its time period is blocked out and becomes unavailable for scheduling lower priority tours and stops.

(page 67) Logsum variables are usually considered the best measures for achieving upward integration. They represent, in upper level utility functions, the expected maximum utility among the alternatives available for lower level choices.

But in model systems this large, formal nested logit models encompassing the entire model system are not practical.

So compromises are made in representing accessibility in the upper level models.

This has been one of the big challenges in implementing activity-based models.

In Sacramento, improving upward integration was an important priority. We used two basic methods to achieve this:

(page 68) Traditionally, logsums that should be calculated across all available times of day ignore that dimension by assuming a most likely time of day.

For example, for work destination choice, mode choice logsums are calculated assuming peak period travel.

(page 69) In Sacramento, we instead used a simulated time of day for each person, using Monte Carlo methods to draw a time of day from a distribution observed among similar cases of travel.

This makes the models that use the logsum appropriately sensitive to policies that are time-of-day specific.

(page 70) It is desirable to use logsums that use the full information of a very specific choice situation. This requires recalculating the logsums for every distinct case in the population.

When the logsums span all available modes and destinations, it becomes too time consuming to do it this way. Therefore, mode-destination logsums are usually not used.

In Sacramento we chose to pre-calculate a large number of mode-destination logsums for each zone, each time the model system is run. For a given zone, we have 84 distinct logsums defined in ways that most affect the logsum values for cases originating in that zone:

--activity purpose; car availability; and proximity of origin parcel to transit.

Then in each case where such a logsum is needed we use the one that matches our case in these dimensions.

(page 71) We use a similar technique to make it possible to use intermediate stop logsums. These measure the attractiveness of making an intermediate stop when the tour origin and destination are given.

(page 72) These kinds of logsums are used to more realistically capture sensitivity to LOS in the upper level models.

(page 73) Now lets look at the model system's equilibration method, which we developed with John Gibb of DKS Associates.

(page 74) Equilibration is important because the demand and assignment must be iterated to achieve consistency in both parts of the system

But performance is an issue. A model system of this size can take a very long time to equilibrate.

(page 75) Here is another picture of the model system, redrawn to focus on equilibration.

--we've shrunk DaySim and the additional trip models to a simple box at the top

--and expanded network assignment to show its internal iterative process.

(page 76) We have an iterative assignment loop, with what I call "a" iterations

--each iteration's assignment result is blended with prior results, in order to speed convergence by reducing or eliminating oscillation around the equilibrium point.

(page 77) This occurs within a larger demand-assignment loop So in the assignment procedure we do a 2-step blending:

--in addition to blending with prior a-iterations to prevent assignment oscillation

--we also blend with prior da-iterations to prevent daoscillation

(page 78) Another technique we use to speed convergence is to run DaySim on only a small fraction of the population in the early iterations

--Each of the first 2 da-iterations simulates less than 1% of the population

--After that, each iteration simulates twice as many persons as the last

--until we finally simulate the entire population on the last iteration

So, in 9 da-iterations, DaySim passes through the population only twice.

(page 79) Last year the SACOG model system was enhanced in two ways to operate more effectively:

--we distributed the day activity schedule model and the traffic assignment to run on multiple processors, to speed up the run time

--we enabled the model to distinguish between long-term and short term effects, as well as to run in a way that we think will be acceptable for FTA New Starts analysis (page 80) For long-term forecasts, DaySim runs all its model components and equilibrates with traffic assignment.

--For short-term effects, the model system is equilibrated in the same way, but comparison scenarios treat the longterm choices modeled in the base scenario as fixed.

--For FTA New Starts, we think it will be okay to do the same thing but also treat the Day Activity pattern and tour destination choice as fixed.

WHAT LIES AHEAD?

(page 81) Now let's turn our attention to what might lie ahead

(page 82) My role in the Denver project was minor, so I will not speak about their model. Tom Rossi is leading that project, so he might have something to say about it.   I will move on to Seattle....

(page 83) ...and tell you about model enhancements of three types that we are preparing to implement in a project that is just beginning.

Seattle is using a model called UrbanSim to simulate economic and land development over time and they want to integrate it with an activity-based model

(page 84) So, instead of simply starting with land use attributes and demographic forecasts

(page 85) The long-term location choices will be moved into the land use simulator,

--which will supply the synthetic population to the activity based model

--after attaching the residential location, usual work locations and usual school locations for each household.

And the activity-based model will be able to provide richer accessibility measures back to the land use model than the current trip-based model system

(page 86) For PSRC we are adding two more models of long term choices that affect people's day-to-day sensitivity to transport policy and travel conditions.

(page 87) And finally we are adding models that represent major aspects of the household's day. One of the criticisms of our model has been that it models the schedules of individual persons, but in real households the members coordinate their activities and travel. In 1994 Moshe and I decided to defer dealing with this explicitly in the models. But since then others have begun to address it, and we will do that also in Seattle.

(page 88) We will model jointly the primary activity for all members of the household.

We will also model half-tours in which one person drops another off at work or school.

And we will model tours in which two or more members of the HH travel together for the entire tour

(page 89) These new features will be challenging enough for our Seattle project.

(page 90) But there are several more enhancements just waiting to be worked on that I believe will be tackled over the next decade.

Humans travel, but it is their vehicles that use energy and pollute. If we can incorporate vehicle type choice and usage as part of the household's activity and travel, then we will be able to more accurately predict the energy and air quality impacts of policies.

(page 91) Our modeling of destination choices essentially ignores the fact that people often park and walk a considerable distance to their activity location. If we include parking location choice in the model system, while accounting for parking supply and the time and cost associated with parking and walking to the desired destination, then we should be able to model better the effects of parking and land use policy in urban centers.

(page 92) For a long time, we have thought that reliability has a big influence on travel choices, but the model systems still aren't incorporating these effects very well.

(page 93) The implementation of activity-based models has so far done nothing to improve the crude ways that are usually used to model commercial traffic.

(page 94) Finally, it looks like the day is approaching when we will be able to implement dynamic traffic assignment models at a regional level and fully integrate them with activitybased models, to provide a comprehensive framework with a consistently high spatial and temporal fidelity.

(page 95) So we definitely have some important work to do for the next several years if we can find agencies willing to fund it.

(page 96) And, although that kind of agency support has only trickled in over the last fifteen years, and the current economic climate is dire, I am cautiously optimistic that the support for activity-based modeling will actually increase somewhat.   (END)